arXiv:1606.05681v1 [cs.CV] 17 Jun 2016

Generating Object Cluster Hierarchies for Benchmarking

Michał Spytkowski[a,1,*], Łukasz P. Olech[a,1], Halina Kwasnicka[a,1]

[a] *Department of Computational Intelligence, Faculty of Computer Science and Management, Wybrzeże Wyspiańskiego 27, Wrocław University of Science and Technology, 50-370 Wrocław, Poland*

Abstract

The field of Machine Learning and the topic of clustering within it is still widely researched. Recently, researchers became interested in a new variant of hierarchical clustering, where hierarchical (partial order) relationships exist not only between clusters but also objects. In this variant of clustering, objects can be assigned not only to leave, but other properties are also defined. Although examples of this approach already exist in literature, the authors have encountered a problem with the analysis and comparison of obtained results. The problem is twofold. Firstly, there is a lack of evaluation methods. Secondly, there is a lack of available benchmark data, at least the authors failed to find them. The aim of this work is to fill the second gap. The main contribution of this paper is a new method of generating hierarchical structures of data. Additionally, the paper includes a theoretical analysis of the generation parameters and their influence on the results. Comprehensive experiments are presented and discussed. The dataset generator and visualiser tools developed are publicly available for use (`http://kio.pwr.edu.pl/?page_id=396`).

Keywords: Hierarchy of Clusters, Benchmark Data Generator, Artificial Data, Tree Structured Stick Breaking Process

*Corresponding author

Email addresses: `michal.spytkowski@pwr.edu.pl` (Michał Spytkowski), `lukasz.olech@pwr.edu.pl` (Łukasz P. Olech), `halina.kwasnicka@pwr.edu.pl` (Halina Kwasnicka)

[1] *URL:* `http://kio.pwr.edu.pl/`

1. Introduction

The enormity of the digital data makes a rapid development of various methods of data analysis, giving an opportunity to analyze the data from a different perspectives, depending on the application. The most popular tasks of data analysis seem to be *regression*, *classification*, and *clustering* [1, 2, 3]. Clustering methods, in general, are still not satisfactory for a variety of data and specific goals, so new methods are still developed.

1.1. Benchmarking in the analysis of clustering methods

Every designer of a new method wants to know how good the proposed method is in comparison with others. Usually such comparison is made by run different methods on some commonly used benchmark datasets and use some of the proposed evaluation measures, e.g., [4, 5, 6, 7, 8, 9, 1]. Such practice is commonly known, here we mention only a very few examples. In [2] the authors used four datasets from the UCI ML repository and few other datasets taken from CHAMELEON, CURE, and BIRCH [10, 11]. Costa et al. [12] used real and synthetic XML data to evaluate the effectiveness and scalability of their method. Authors of [13] evaluated the performance of three different clustering methods according to four indices. Experiments were made using artificial and real-life datasets with the different number of clusters ranging from two to ten. Mirzaei and Mohammad [14] proposed an ensemble of hierarchical clustering, where obtained dendrograms are converted to matrices, which are further aggregated into a final matrix – a basis for forming the final clustering. The combination of hierarchical clustering was compared to the combination of partitional clustering. The experiments were made using twelve small benchmark datasets (from 98 to 353 instances) and twelve large datasets – from 625 to 5473 instances. Summing up, the evaluation of clustering techniques is made experimentally, using some developed quality measures that take into account some properties of the obtained clusters. The methods are run using some different datasets to collect quality measure values on different input data. This approach

is common for various Machine Learning methods [4, 5, 6, 7, 8, 9, 15, 16, 17]. In this paper we focus on a special type of clustering, a particular approach to hierarchical clustering, namely – a generation of a *object cluster hierarchy*. Hierarchical clustering paradigm [2, 12, 15, 18], agglomerative or divisive approach, produces a dendrogram showing all levels of agglomerations. Although there is a hierarchy relation between clusters, there is **no hierarchy relationship between objects**. It is because all grouped objects are assigned only to the leaves and clusters are the result of cutting the tree at different levels. Any node in the tree, except leaves, does not have assigned objects. The structure of generated clusters is flat.

1.2. Brief introduction to Object Cluster Hierarchies

Our research concentrates on another **paradigm that extends hierarchical cluster analysis and allows assigning objects to each node** in the produced hierarchy tree [19, 20, 21]. We are developing a method which allows objects to be assigned to any node of the hierarchy tree, what is enough to receive a new hierarchy relation that exists between objects. Beside that we formulate three important requirements [20]:

- *Inheritance* – every object belonging to a given group also belongs to the parents groups, up to the root;
- *Retention* – objects do not need to be located in the tree's leaves;
- *Variance* – groups located lower in the hierarchy are more specific, i.e., every child group has to have not higher variation than its parents.

One can put the question: why a generation of hierarchical structures reflecting hierarchies between objects is worth to study? One of the first work on this subject is [19] where the authors pointed out that many data arise from a latent hierarchy, for example, a set of text documents or images. Such data can be modeled by a hierarchical structure, one possible way to discover the unobservable structure is to infer it during a learning process. Adams et al.

in [19] proposed the very flexible, nonparametric method allowing to develop trees of unbounded width and depth. The objects can be assigned to any node in a created tree. Also, some nodes can be empty, i.e., without assigned objects. Above properties cause that the method is very interesting however we are not satisfied with its properties: objects belonging to the child node (child group) can vary more than objects assigned to the parent' nodes. Starting from this method as the prototype, we propose a new version of it, mainly but not only, by changing the kernel. After developing the authors' methods, we met the troubles with comparing our method with others, especially with our inspiration [19] and Bayesian Rose Trees [22]. The problem is twofold: firstly, there are no appropriate evaluation measures dedicated to such hierarchies of objects, they should reflect the desired properties listed above. Secondly – there is no available benchmark datasets with known properties, which could be used for testing different methods of hierarchical structures generation. Currently we try to overcome both difficulties. The first problem is a lack of proper measures dedicated to hierarchies with assumed requirements. The partial solution to this problem is presented in [23]; however, we have developed more new, external and internal validation techniques which we plan to publish soon. Their experimental study requires appropriate datasets. So we decided to develop the method of generation the data structures with known characteristics, useful for generation of benchmark datasets that can be used by every researcher who is involved in the subject.

1.3. Contribution and problem description

The aim of our research is to **fill the gap in the presence of commonly accessible benchmark datasets**. The contribution of the paper is strictly scientific – we have developed a **new method generating the hierarchical structures of data with assumed properties**. The additional benefit for researchers can be a number of new benchmarks – hierarchical structures of data. Making our method freely available from the website (together with the instruction of how to use it) we allow researchers to generate new structures

of data according to their needs. This will allow analyzing the quality of the evaluated methods – how generated structures are adequate to the input data, and to compare the quality of various methods of generating hierarchical structures, with the assumption that the used indices reflect required properties of the methods.

The design of benchmark data is not a trivial task [24]. Checking particular properties of the method requires the appropriate input data, revealing the desirable feature of the tested method, or the characteristics of the real data that the method is going to work with. Conclusions obtained on specific datasets often cannot be applied to other datasets. However, conclusions based on data characteristics can be easily generalized what allows predicting the performance of evaluated method [25]. Chen in his thesis [25] controls shape, volume, and size of particular classes of generated datasets. The control by the parameters value of the generative models is desirable because it is convenient and facilitates the comprehensive evaluation process. A new benchmark dataset with production methodology is presented by O'shea et al. in [26]. The benchmark data concerns short text semantic similarity algorithms. The authors underline importance of benchmark datasets in the Machine Learning area.

Developing our method, we have taken into account all aspects of the potential use of the method as well as the generated datasets. The proposed tool (and datasets) can be used to test methods of clustering that should endure several properties of output models, they correspond to the three requirements, mentioned earlier:

- the data can be assigned to any node in the hierarchy, not only leaves,
- the standard deviation of data should always be lower (at most equal) in a child node than in the corresponding parent node,
- the data in a node should be more similar to data in parent and child clusters than to unrelated clusters located in different subtrees in the hierarchy.

In this paper we present details of the model of generated structures. The implemented tool (available for everyone) was used to study the influence of the model parameters on the characteristic of generated datasets. We analytically evaluated the average probability and the standard deviation of the data (objects) staying at particular levels of the hierarchy tree, depending on the parameters of the model. We also present some features of our tool by analysing the generated data. The tools together with example generated datasets are free, and they are accessible from `http://kio.pwr.edu.pl/?page_id=396`.

The paper is organized as follows. The next section presents the model and the generation process. In Section 3 we describe the meaning of used parameters. These parameters control such features of generated structures as the hierarchy deep and wide, a distribution of data crossing the levels of the hierarchy, etc. Conducted experiments and their results are shown and discussed in Section 4. We verified the correctness of analytical and intuitive properties of the generator and how the different parameters affect the outcome of generation. The last section concludes the paper.

2. Generator model

The data generated by this tool can be seen as data coming from an infinite mixture model. Most commonly an infinite mixture model is composed of infinite, indexed distributions from which the data is drawn. In such a case the mixture weights can be drawn from a Dirichlet distribution, using the Stick Breaking Process. Generally, the distributions for the mixture components are unrelated to each other. This generator uses a similar approach, however it bases its mixture weights on the Tree Structured Stick Breaking Process, as described in [19], which arranges the mixture components into a hierarchical structure. This structure also defines the relationship between the mixture components, that is – parameters of the child distribution are based on the parent's parameters.

Hierarchies generated using the Stick Breaking Process posses the following

characteristics:

1. Every node potentially has an infinite number of children;
2. The children of a node are indexed and ordered, however this indexing is not important after the generation process finishes;
3. Every node in the tree can have child nodes, the potential depth of the tree is infinite;
4. The hierarchy may contain empty nodes, that is, nodes that do not generate any data;
5. The child node parameters are generated based on the parent parameters and a kernel describing the transition;
6. The shape of the generated tree is dependent on a number of control hyperparameters.

Throughout this paper following symbols will be used to describe the generator and the generated model:

X - set of all data point, or objects,

x_i - an object or data point with unique identifier i represented by a vector of features,

Θ - set of all clusters,

ϵ - specific cluster from Θ,

ϵ_{x_i} - cluster of object x_i,

$\epsilon\epsilon_i$ - in cases where a hierarchy of clusters exists $\epsilon\epsilon_i$ is the $i-th$ child of cluster ϵ,

$\epsilon_\varnothing$ the root cluster of the cluster hierarchy, if one exists,

X_ϵ - set of all objects in cluster ϵ,

X_{E_ϵ} - set of all objects in hierarchy subtree beginning with node ϵ,

$|S|$ - cardinality of set S,

$|\epsilon|$ - depth of node ϵ.

Additionally to the symbols above several values are provided for the generator as parameters:

d - the number of dimensions for the data,

n - the number of data to be generated,

$\alpha(\epsilon)$ - one of the control parameters for tree structure, described in depth later, we assume $\alpha(\epsilon) = \alpha_0 \lambda^{|\epsilon|}$,

γ - second control paramter for tree structure,

p, q - parameters for the kernel,

$\theta_\varnothing$ - distribution for the root node.

Having these parameters we can draw conditional probabilities used to determine which node data is generated from. The first is the conditional probability of a datum remaining in node ϵ, at depth $|\epsilon|$, when entering the node:

$$\nu_\epsilon = P(x \in X_\epsilon | x \in X_{E_\epsilon}) \tag{1}$$

$$\nu_\epsilon \sim Beta(1, \alpha(\epsilon)), \quad \alpha(\epsilon) = \alpha_0 \lambda^{|\epsilon|} \tag{2}$$

The second is the conditional probability of a datum traveling down subtree $\epsilon\epsilon_i$ if it does not remain in node ϵ and does not travel down sibling subtrees with a lower index ($\epsilon\epsilon_j, j < i$):

$$\psi_{\epsilon\epsilon_i} = P(x \in X_{E_{\epsilon\epsilon_i}} | x \in X_{E_\epsilon} \wedge x \in X_\epsilon \wedge \neg\exists_{j<i} x \in X_{E_{\epsilon\epsilon_j}}) \tag{3}$$

$$\psi_{\epsilon\epsilon_i} \sim Beta(1, \gamma) \tag{4}$$

Additionally we need to define the kernel. We begin with a defined root node distribution given as a starting parameter of the generation method. The values defined at this point are the means and standard deviations for each of the Gauss distributions in the d different dimensions:

$$\theta_\varnothing = (Gauss(\mu_{\varnothing 1}, \sigma_{\varnothing 1}), ..., Gauss(\mu_{\varnothing d}, \sigma_{\varnothing d})) \tag{5}$$

Then from there, for any node for which we need the distribution we can draw the distribution based on the parent's distribution. The child's mean values are drawn directly from the parent distribution and the child's standard deviation based on a scaling factor ($\Delta\sigma_n$) drawn from a $Beta$ distribution. The values

are taken separately for each dimension:

$$\theta_{\epsilon\epsilon_i} = (Gauss(\Delta\mu_1, \sigma_{\epsilon 1}\Delta\sigma_1), ..., Gauss(\Delta\mu_d, \sigma_{\epsilon d}\Delta\sigma_d)) \tag{6}$$

$$\Delta\mu_n \sim Gauss(\mu_{\epsilon n}, \sigma_{\epsilon n}), \quad n = 1...d \tag{7}$$

$$\Delta\sigma_n \sim Beta(p, q), \quad n = 1...d \tag{8}$$

With the kernel defined we can now generate data from the model. We begin with the hyperparameters and the probability distribution for the root node $(\theta_{\epsilon_\varnothing} = (Gauss(\mu_{\epsilon_\varnothing 1}, \sigma_{\epsilon_\varnothing 1}), ..., Gauss(\mu_{\epsilon_\varnothing d}, \sigma_{\epsilon_\varnothing d})))$.

The following process continues until we generate n points:

Step 1: If $|X| < n$ go to **Step2**, else **end**.

Step 2: Randomly draw an insertion point $i_x \sim Uni(0, 1)$, $i_x \in (0, 1)$.

Step 3: Set the root node as the current node ($\epsilon_c := \epsilon_0$), depth is 0 ($|\epsilon_c| := 0$).

Step 4: If ν for the current node is not yet known, draw the value. $\nu_{\epsilon_c} \sim Beta(1, \alpha_0\lambda^{|\epsilon_c|})$.

Step 5: If $i_x \leq \nu_{\epsilon_c}$ then $x \sim \theta_{\epsilon_c}$, the point belongs to the current node ($X_{\epsilon_c} := \{x\} \cup X_{\epsilon_c}$), go to **Step 1**, else move on to **Step 6**.

Step 6: Adjust i_x to new value: $i_x := (i_x - \nu_{\epsilon_c})/(1 - \nu_{\epsilon_c})$.

Step 7: Set current child node index ($\epsilon_c\epsilon_i$) to the first child node of the current node: $i := 0$.

Step 8: If ψ for the current child node is not yet known, draw the value: $\psi_{\epsilon_c\epsilon_i} \sim Beta(1, \gamma)$.

Step 9: If θ for the current child node is not yet known, draw the values based on the parent of the node:

$\theta_{\epsilon_c\epsilon_i} = (Gauss(\Delta\mu_1, \sigma_{\epsilon_c 1}\Delta\sigma_1), ..., Gauss(\Delta\mu_d, \sigma_{\epsilon_c d}\Delta\sigma_d))$,

$\Delta\mu_1$ is drawn from first dimension of parent node ($\Delta\mu_n \sim Gauss(\mu_{\epsilon_c 1}, \sigma_{\epsilon_c 1})$),

...

$\Delta\mu_d$ is drawn from d-th dimension of parent node ($\Delta\mu_d \sim Gauss(\mu_{\epsilon_c d}, \sigma_{\epsilon_c d})$),

$\Delta\sigma_1 \sim Beta(p, q)$,

...

$\Delta\sigma_d \sim Beta(p, q)$.

Step 10: If $i_x \leq \psi_{\epsilon_c\epsilon_i}$ go to **Step 11**, else go to **Step 12**.

Step 11: Adjust the value of i_x to new value: $i_x := i_x/\psi_{\epsilon_c\epsilon_i}$. Make the current child the current node ($\epsilon_c := \epsilon_c\epsilon_i$) and increase depth ($|\epsilon_c| := |\epsilon_c| + 1$). Go to **Step 4**.

Step 12: Adjust the value of i_x to new value $i_x := (i_x - \psi_{\epsilon_c\epsilon_i})/(1 - \psi_{\epsilon_c\epsilon_i})$. Increment child index of currently relevant child node ($i := i + 1$). Go to **Step 8**.

For better understanding, the generation process described above is illustrated in block diagram Figure 1.

3. Parameter selection

Different applications will require the generation of different benchmark data to best suit the desired outcome. The primary concern is the structure of the hierarchy, that is if the hierarchy is high or short, wide or narrow, as well as the distribution of data cross the levels of the hierarchy and also specific nodes. Additionally the difference between data in parent and child nodes can also be important in some cases. All of them are controlled by a number of parameters in the model:

1. hierarchy depth: α_0, λ or in a more general sense – the $\alpha(\epsilon)$ function,
2. hierarchy width: γ,
3. data specificity: p, q,
4. starting node distribution: $\theta_\varnothing$.

3.1. Controlling the hierarchy depth

The depth of the hierarchy is controlled by the α. The higher the probabilities of data remaining in nodes, the less data will travel deep down the tree and thus the tree will be shallower. On the other hand, if the probability is lower the data will, on average, travel deeper into the tree before stopping. The average probability of data remaining in a given node is based on the selected

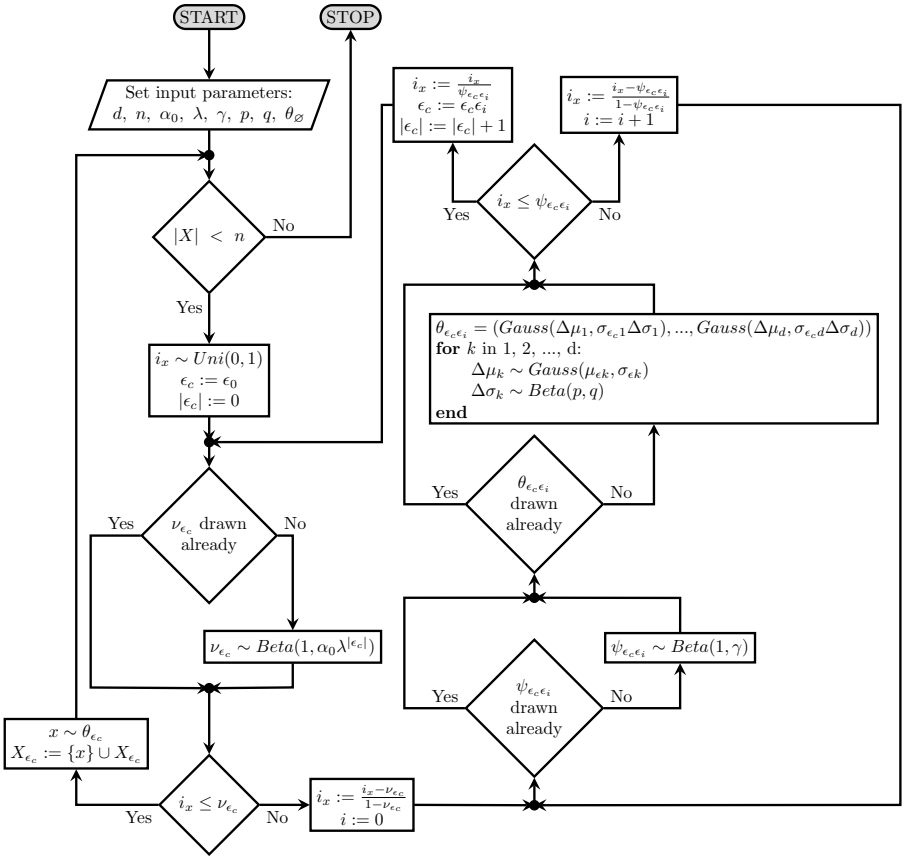

Figure 1: Block diagram of Object Cluster Hierarchy generator

α function which in turn is the basis for estimating the behaviour of the tree based on the α_0 and λ parameters:

$$\mathbf{E}[x \in X_\epsilon, |\epsilon| = 0] = \frac{1}{1 + \alpha_0} \tag{9}$$

$$\mathbf{E}[x \in X_\epsilon, |\epsilon| = n] = \frac{\prod_{i=0}^{n-1} \alpha_0 \lambda^i}{\prod_{j=0}^{n} (1 + \alpha_0 \lambda^j)} \tag{10}$$

Additionally the variance can also be calculated:

$$\mathbf{var}[x \in X_\epsilon, |\epsilon| = 0] = \frac{\alpha_0}{(1 + \alpha_0)^2 (2 + \alpha_0)} \tag{11}$$

$$\mathbf{var}[x \in X_\epsilon, |\epsilon| = n] = \frac{2\prod_{i=0}^{n-1} \alpha_0 \lambda^i}{(1+\alpha_0\lambda^n)\prod_{j=0}^{n}(2+\alpha_0\lambda^j)} - (\mathbf{E}[x \in X_\epsilon, |\epsilon| = n])^2 \tag{12}$$

Taking the two above value sets into account, as well as the data generation procedure, it is possible to estimate the shape of the tree. As parameters are continuous so is the spectrum of likely outcomes:

- $\alpha_0 = 1, \lambda = 1$: The structure of the tree is chaotic and hard to predict, the further away the parameters move from these values the more stable the tree becomes;
- $\alpha_0 < 1, \lambda \leq 1$: Shallow structure, data located primarily at the top of the tree;
- $\alpha_0 < 1, \lambda > 1$: Similar to the above case, depth of the tree increases but most data is located at the top of the tree;
- $\alpha_0 > 1, \lambda < 1$: The structure is deep but data is not located at the top, the bigger α_0 starts out and smaller λ is the more data will move down the tree into the central or lower region;
- $\alpha_0 > 1, \lambda \geq 1$: Deep structure, data located primarily at top of the tree but spread out.

3.2. Controlling the hierarchy width

The width of the tree is based on the value of the γ parameter. Assuming that $x \in X_{E_\epsilon}$ and $x \notin X_\epsilon$ the average probability of data being generated from a specific subtree (based on index) can be used to estimate the number of children a node can potentially have:

$$\mathbf{E}[x \in X_{E_\epsilon \epsilon_i}, i = 1] = \frac{1}{1+\gamma} \tag{13}$$

$$\mathbf{E}[x \in X_{E_\epsilon \epsilon_i}, i = n] = \frac{\gamma^{n-1}}{(1+\alpha_0\lambda^j)^n} \tag{14}$$

Variance for these values can also be calculated:

$$\mathbf{var}[x \in X_{E_\epsilon \epsilon_i}, i = 1] = \frac{\gamma}{(1+\gamma)^2(2+\gamma)} \tag{15}$$

$$\mathbf{var}[x \in X_{E_\epsilon \epsilon_i}, i = n] = \frac{2\gamma^{n-1}}{(1+\gamma)(2+\gamma)^n} - (\mathbf{E}[x \in X_{E_\epsilon \epsilon_i}, i = n])^2 \tag{16}$$

Once more, taking the generation procedure into account and the influence of the parameter on the estimated values to judge the general tendencies of the model:

- $\gamma = 1$: The number of children is chaotic and difficult to judge;
- $\gamma < 1$: Narrower tree, less children per node on average;
- $\gamma > 1$: Wider tree, more children per node on average.

3.3. Controlling data specificity

When a new group is considered, the parameters for that group's distribution are drawn based on the parent distribution and the p and q kernel parameters. An important aspect of the generated model is that data becomes more specific in lower nodes. This behaviour is always present, however the values taken for the kernel change the average proportion of data standard deviation between the parent and child. This is based on the expected standard deviation of the new node compared to the old node (taken separately in each dimension):

$$\mathbf{E}[\sigma_{\epsilon\epsilon_i d}] = \sigma_{\epsilon d}\frac{p}{p+q} \tag{17}$$

$$\mathbf{var}[\sigma_{\epsilon\epsilon_i d}] = \sigma_{\epsilon d}\frac{pq}{(p+q)^2(p+q+1)} \tag{18}$$

By selecting p and q the rate at which the nodes become more specific can be altered. The lower the mean is the more specific every child will be (on average), the higher the variance is the more variety there will be in how the child nodes relate to their parent.

3.4. Influence of starting distribution on results

Due to the relative nature of the model (i.e., the specific values generated from the model are calculated relative to each other, starting from the set root distribution) the choice of initial distribution parameters is not important. The data generated from the model can be scaled afterwards to any desired values as

well as moved in any direction along any dimension. Due to this the generator assumes a value for $\theta_{\varnothing}$:

$$\theta_{\varnothing} = (Gauss(0, \sigma_{max}), ..., Gauss(0, \sigma_{max})) \tag{19}$$

Data generated from the model can be then post processed to a more desirable spread of values. This is done by applying scaling and translation to all data generated by the model as well as the parameters of each group node.

3.5. Post-processing

The data also undergoes one form of post-processing after being generated. This process, referred in this paper as *reassignment*, moves the data between clusters in such a way that each object belongs to the cluster it is most likely to be generated by:

$$\forall_{x \in X} \left(x \in X_{\epsilon_a} \Leftrightarrow !\exists_{\epsilon_b \neq \epsilon_a} L(x|\theta_{\epsilon_a}) < L(x|\theta_{\epsilon_b}) \right) \tag{20}$$

The process does not modify the clusters number, relations or parameters in any way. It merely relocates data in order to reduce noise and produce cleaner clusters.

4. Experiments

Tests performed on the generator have a number of different goals. Primarily the tests served to investigate the correctness of analytical and intuitive properties of the generator parameters. Thus, a large part of the experiments serves to produce data to be investigated with the purpose of verifying how the different parameters affect the outcome of generation. Secondly, the experiments serve to visualise various properties of the generated hierarchies as a reference for anyone interested in using the generator. The presented values give a clearer view of the results depending on the various parameters, helping to judge the best parameter set for a given use. Thirdly, the results are taken separately for the raw datasets as generated from the statistical model, and the reassigned datasets.

The goal of the reassignment process is reducing noise in the generated data by moving objects to the node for which the likelihood of being drawn from that node is the highest. The purpose of these measurements was determining the influence of reassignment on the overall structure of the hierarchy.

Data presented in this section was obtained by averaging over 100 generations for each of the parameter sets. Each of the parameter sets represents a different hierarchy structure, some parameters remain constant across all experiments Table 1. The experiments were then repeated twice for each generated data set: once using the initial assignment of data to nodes and a second time after a reassignment of data. A number of quantitative measures were used to

Set	α_0	λ	γ
s00	1	0,5	0,2
s01	1	1,0	0,2
s02	1	1,0	1,0
s03	5	0,5	0,2
s04	5	1,0	0,2
s05	5	0,5	1,0
s06	25	0,5	0,2
s07	25	0,5	1,0

Table 1: Generator's parameters used to create experimental data set from *s00* to *s07*. The remaining parameters are shared between test sets: $n = 10000$, $d = 2$, $p = 1$, $q = 5$, $\sigma_{min} = 0.05$ and $\sigma_{max} = 10$.

investigate the properties of the generated hierarchies:

- N – number of nodes in hierarchy, averaged over all hierarchies generated,
- L – number of leaves in hierarchy (nodes with no children or only empty children), averaged over all hierarchies generated,
- D – depth of the hierarchy, averaged over all hierarchies generated,

- B – breadth of the hierarchy, averaged over all levels in a hierarchy, then over all hierarchies generated,
- P – the average length of all paths in a hierarchy, averaged over all hierarchies generated.

Since the two last values (B and P) are averages of averages, the value given to them is not the standard deviation but rather the average standard deviation over all generated hierarchies. All defined measures are reported separately for the regular and reassigned hierarchies in Tables 2 and 3.

Set	*Nodes*		*Leaves*		*Depth*		*Breadth*		*Path length*	
	$\bar{N}$	$\sigma_{\bar{N}}$	$\bar{L}$	$\sigma_{\bar{L}}$	$\bar{D}$	$\sigma_{\bar{D}}$	$\bar{B}$	$\bar{\sigma}_{\bar{B}}$	$\bar{P}$	$\bar{\sigma}_{\bar{P}}$
s00	17,58	7,80	8,75	4,33	4,06	0,82	3,40	2,14	2,86	0,93
s01	95,23	53,06	31,00	17,85	11,73	2,80	7,14	5,10	6,30	2,50
s02	556,81	329,74	271,80	188,79	12,33	2,20	40,44	41,33	5,40	1,99
s03	58,19	21,33	25,84	10,62	6,41	0,73	7,84	5,97	4,36	1,13
s04	3090,88	944,13	483,62	187,81	52,54	6,83	58,14	59,92	19,39	7,88
s05	485,43	149,83	297,62	108,54	6,88	0,55	61,69	64,09	4,25	1,00
s06	175,71	61,50	67,57	26,01	8,83	0,64	17,84	14,97	6,20	1,32
s07	2071,07	536,50	1109,17	367,70	9,27	0,51	201,88	223,79	5,88	1,16

Table 2: Accumulative characteristics of generated hierarchies **without** the reassignment procedure. Average $\bar{X}$ values together with standard deviation $\sigma_{\bar{X}}$ (or average of standard deviations $\bar{\sigma}_{\bar{X}}$) are provided.

Set	*Nodes*		*Leaves*		*Depth*		*Breadth*		*Path length*	
	$\bar{N}$	$\sigma_{\bar{N}}$	$\bar{L}$	$\sigma_{\bar{L}}$	$\bar{D}$	$\sigma_{\bar{D}}$	$\bar{B}$	$\bar{\sigma}_{\bar{B}}$	$\bar{P}$	$\bar{\sigma}_{\bar{P}}$
s00r	18,11	8,17	9,30	4,68	4,05	0,82	3,50	2,23	2,86	0,92
s01r	98,56	54,99	34,76	19,50	11,70	2,82	7,40	5,31	6,30	2,48
s02r	642,02	367,21	363,20	211,42	12,23	2,19	46,88	47,62	5,43	1,96
s03r	59,88	22,11	27,64	11,33	6,40	0,74	8,07	6,15	4,35	1,14
s04r	3099,03	936,39	597,36	176,93	52,05	6,96	58,90	60,44	19,17	7,78
s05r	552,61	151,38	366,31	100,87	6,87	0,56	70,30	72,92	4,25	1,03
s06r	180,47	63,97	73,14	28,05	8,81	0,65	18,34	15,40	6,16	1,35
s07r	2310,21	524,57	1375,60	302,29	9,27	0,51	225,19	250,42	5,87	1,18

Table 3: Accumulative characteristics of generated hierarchies **with** the reassignment procedure. Average $\bar{X}$ values together with standard deviation $\sigma_{\bar{X}}$ (or average of standard deviations $\bar{\sigma}_{\bar{X}}$) are provided.

The rest of the presented experiments figs. 2 to 9 consist of histograms averaged over the 100 generated hierarchies for each parameter set:

- average number of objects per node per level,
- average width per level,
- average number of leaves per level,
- average number of children per node per level
- average number of nodes with a given number of children

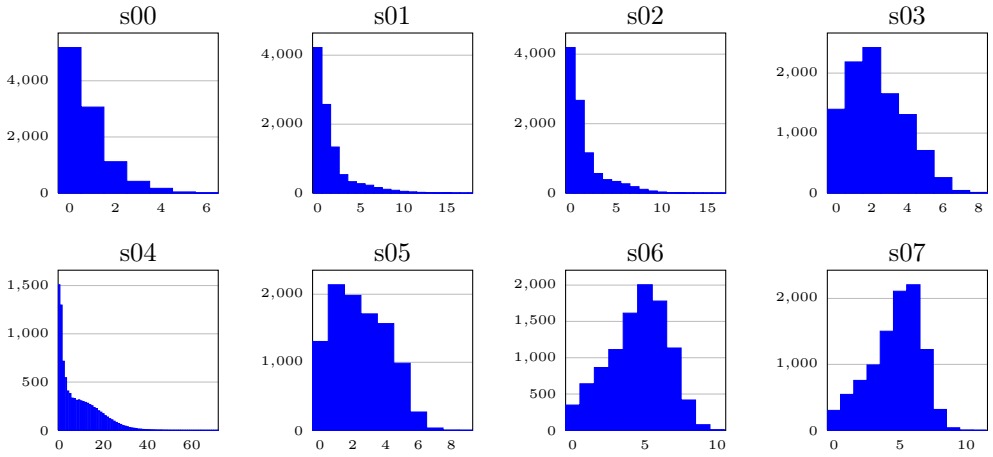

Figure 2: Average distribution of data instances among hierarchy levels **without** execution of reassignment procedure. Vertical axes show the number of instances and horizontal axes indicate hierarchy level.

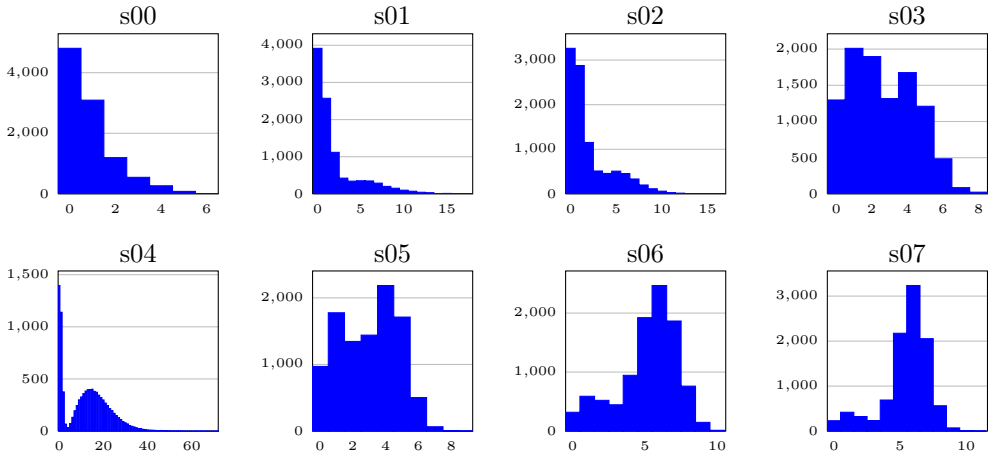

Figure 3: Average distribution of data instances among hierarchy levels **with** execution of reassignment procedure. Vertical axes show the number of instances and horizontal axes indicate hierarchy level.

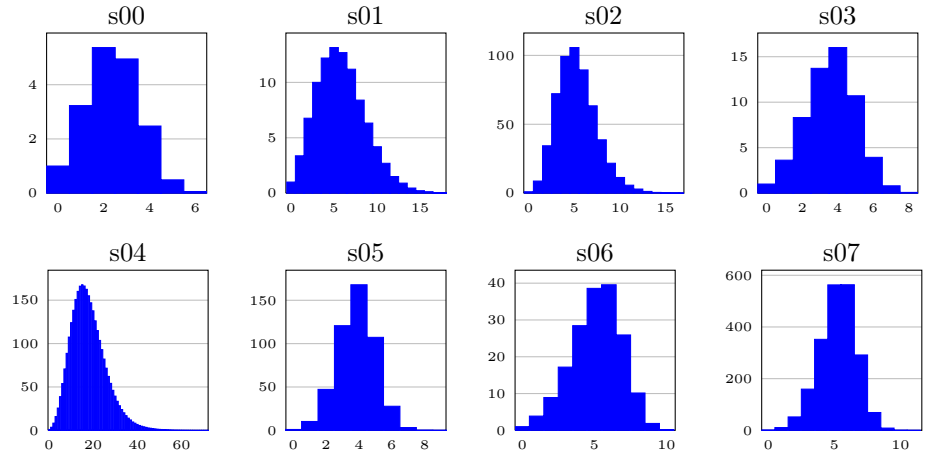

Figure 4: Average hierarchy width (B) on every hierarchy level (number of nodes on every level) **without** execution of reassignment procedure. Vertical axes show hierarchy width and horizontal axes indicate hierarchy level.

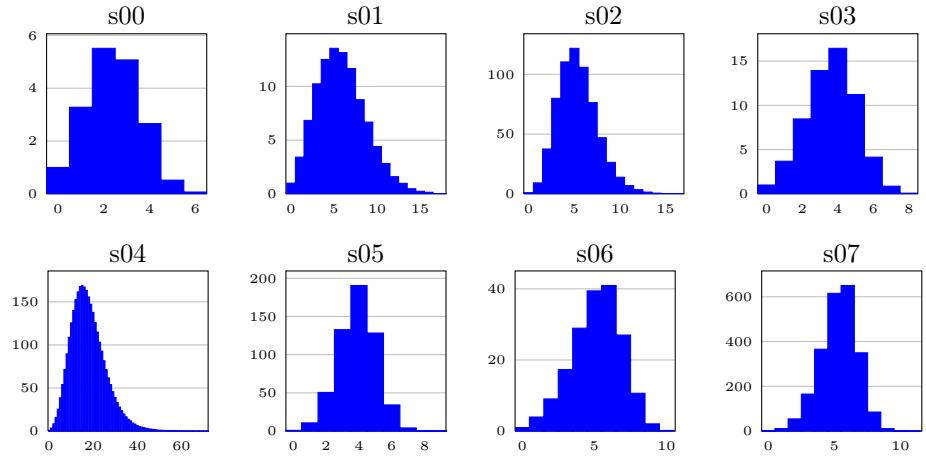

Figure 5: Average hierarchy width (B) on every hierarchy level (number of nodes on every level) **with** execution of reassignment procedure. Vertical axes show hierarchy width and horizontal axes indicate hierarchy level.

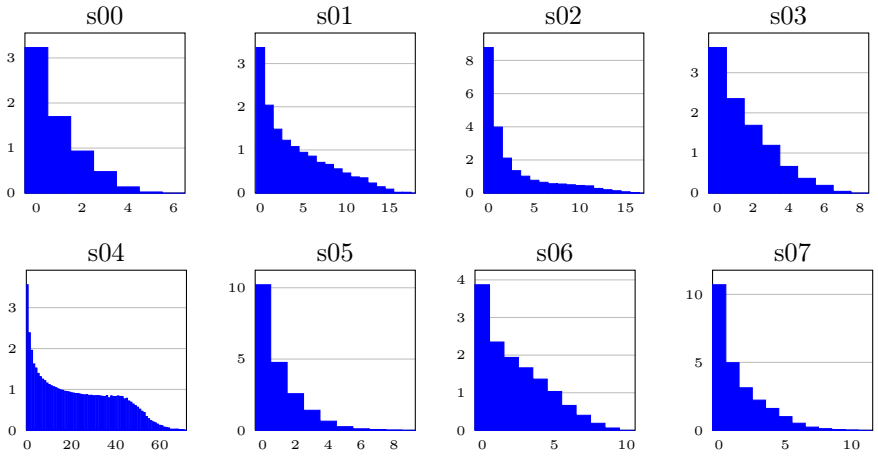

Figure 6: Distribution of the average number of children per node among hierarchy levels **without** execution of reassignment procedure. Vertical axes show the number of children and horizontal axes indicate hierarchy level.

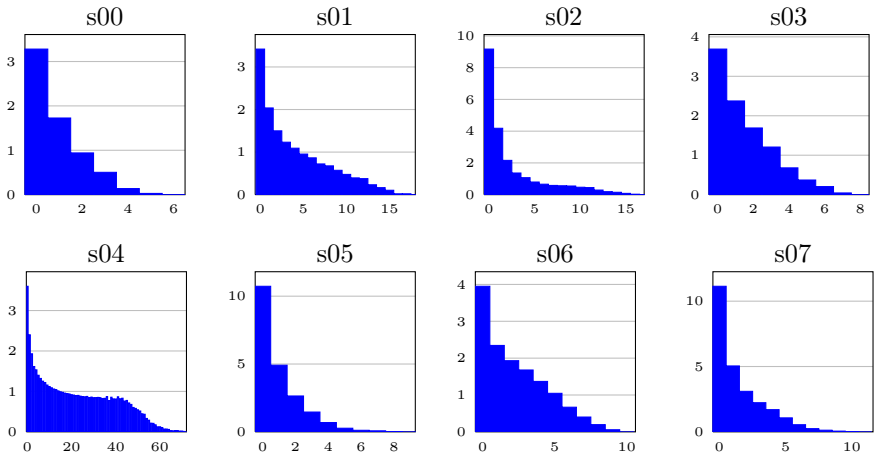

Figure 7: Distribution of the average number of children per node among hierarchy levels **with** execution of reassignment procedure. Vertical axes show the number of children and horizontal axes indicate hierarchy level.

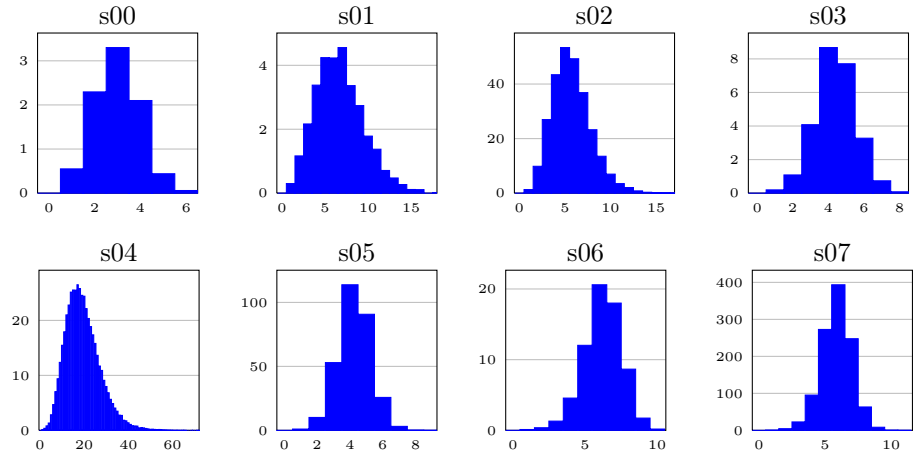

Figure 8: Average number of leaf nodes (L) on every hierarchy level **without** execution of reassignment procedure. Vertical axes show the number of children and horizontal axes indicate hierarchy level.

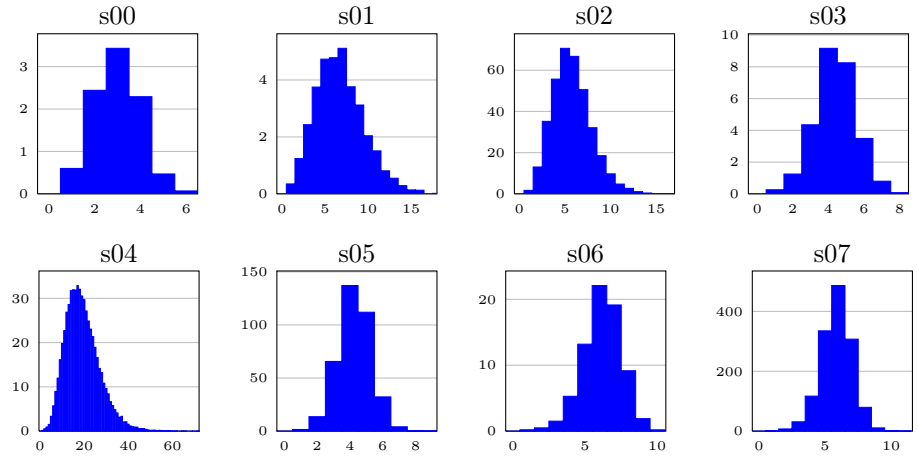

Figure 9: Average number of leaf nodes (L) on every hierarchy level **with** execution of reassignment procedure. Vertical axes show the number of children and horizontal axes indicate hierarchy level.

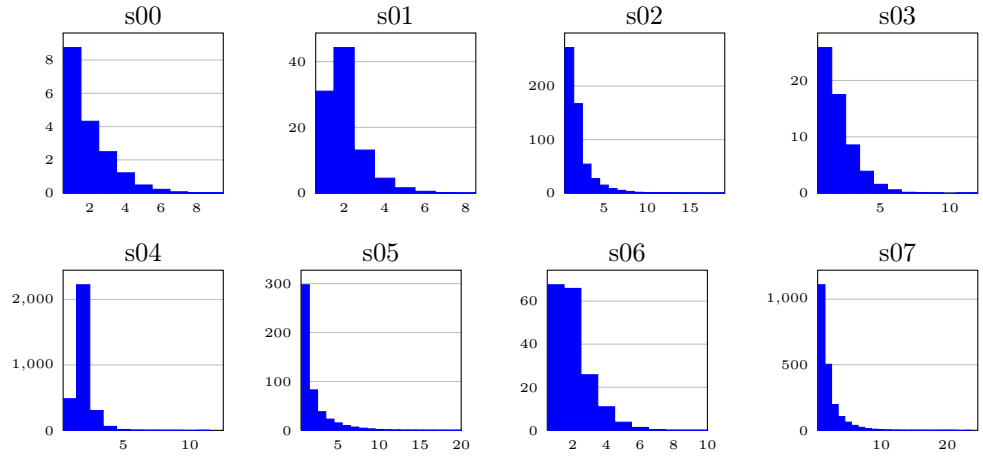

Figure 10: Average number of child nodes for every node in generated hierarchies **without** execution of reassignment procedure. Horizontal axes show the number of children and vertical axes show the number of occurrences (count) in the hierarchies.

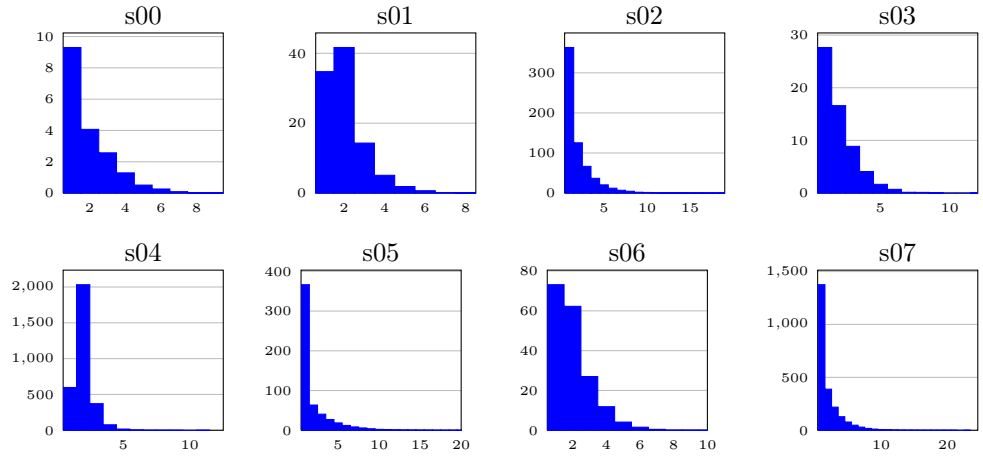

Figure 11: Average number of child nodes for every node in generated hierarchies **with** execution of reassignment procedure. Horizontal axes show the number of children and vertical axes show the number of occurrences (count) in the hierarchies.

The data accumulated from the generated hierarchies can be confronted with prior analytical estimations of the effect that parameters have on the structure

of the hierarchy. The simpler issue is the γ parameter. This parameter is responsible for the formation of children and as such the breadth of the hierarchy. We can see that for datasets that differ only by the γ value ($s01$ and $s02$ or $s06$ and $s07$), the distribution of data per level is very similar (Figures 4 and 5) – it would be controlled by the α function, which does not change in this case. On the other hand, there is a significant change in the width of the hierarchy, approximately by one order of magnitude (10 times higher for higher γ), as is predicted by the prior analysis.

The influence of the α_0 and λ are more difficult to describe concisely, as the two parameters are interwoven together within the α function, Equation (2). However the most significant presentation of the influence of this function is presented by comparison of the $s00$, $s01$, $s02$ and $s04$ dataset when compared with the $s03$, $s05$, $s06$ and $s07$ datasets. The first set of hierarchies has a clear tendency to retain data at higher levels more so than at lower ones (Figures 2 and 3). In comparison, the other four hierarchies have the main mass of data lower than at the root level. Especially with the final two datasets ($s06$ and $s07$) the majority of objects are located close to the 5th level of the hierarchy. For these two datasets we can see that α starts out as a higher value and declines at lower levels. For a high value of α the probability of retaining data in a node is low on average (compare the influence of parameters on the Beta distribution). Thus the higher levels do not retain data, but as the value of α drops over time more of it gathers in the lower levels of the hierarchy before finally the remaining data is passed on to the lowest levels. From this we come to an important conclusion about the importance of these two parameters for applications of the generator. In cases where it is undesirable to have many generic (root level) objects and it is important to have clearly distinct specific (lower level) objects, these two parameters must have values similar to those of $s06$ and $s07$ – high starting α_0 and a λ that controls the decline of the function value over levels. This behaviour was earlier predicted from the analytical study of the parameters also, see Equations (17) and (18), however the data shows the behaviour more clearly. It appears that the critical point for the parameters

is the level at which the value of α drops below 1, though this was not fully explored in this paper and is left for further investigation.

A very prominent behaviour of the generator seen in all test cases is producing what will be referred to from this point onward as trailing divisions of data. Trailing divisions occur when the generator attempts to split small remaining partitions of data. This occurs both in right (higher index) children of a populated node and lowers down the hierarchy (Figure 12). In both cases it is possible to observe large numbers of small nodes (i.e., with a low number of children, usually one or zero – that is to say the node is a leaf node), as well as many nodes that are not populated with data. Trailing divisions reveal the fractal nature of trees generated by the procedure, which manifests itself both when producing direct children for a node (horizontal self-similarity) and going down the hierarchy (vertical self-similarity).

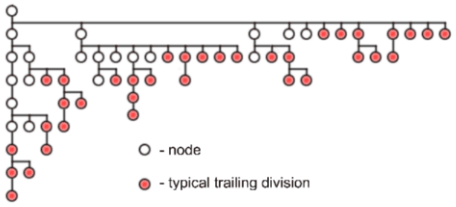

Figure 12: A simple schema of the location of training divisions.

Horizontally the above phenomenon can be visualised as the ordered set of all children of any node being statistically similar to the ordered set of all children of the node except the first one. It is a direct effect of the Tree Structured Stick Breaking process [19]. Similarly to the above, from a vertical point of view, in any tree where $\lambda = 1$ all sub-trees of a node are statistically similar to that node. In the presented experimental data, these trailing divisions are visible as the falling off "tail" in the data per level histograms as well as the cause of the high deviation of width within the tree. Unfortunately, due to the nature of the TSSB distribution, it is impossible to avoid this behaviour without post-processing. No forms of such post-processing were employed for the experiments

presented in this paper.

As an additional observation to the above, the γ value does not change by level. Because of this, a very important factor in considering how many children a node will have is instead the number of data passing through the node. Bigger nodes that are located higher up are more likely to have more children than smaller nodes lower down the hierarchy. Additionally, the more children a node has, the more of them will be small nodes, i.e., nodes through which few objects pass, resulting primarily in leaf nodes or nodes with single children. This behaviour of smaller nodes also transfers to lower down the hierarchy where less data reaches leading to a similar behaviour.

Finally, the reassigned test cases show a tendency for data to move down hierarchy levels. Intuitively the groups located lower down in the hierarchy are more specific and potentially conflicting data would be prone to moving down into the more specific child clusters. However despite this tendency the hierarchies retain most of their previous characteristics, simply with the mass of data being shifted downward towards the lower levels of the hierarchy. Due to the post-processing applied to these datasets, they can be better suited for initial testing of grouping methods. In most cases testing using both types of datasets (unfiltered and filtered) may be the preferred and most valuable approach in every case where the features of the objects are considered.

5. Conclusions

The experiments presented in the previous section serve to highlight both the strengths and weaknesses of the proposed generator. A prominent strength is a high range of different tree structures that can possibly be generated and the ability to control these structures using the parameters. From a practical point of view it is also important that the parameters can be separated into groups, each controlling a different aspect of the hierarchy (α_0, λ – vertical distribution of data, γ – width, p, q – specificity) allowing for fine tuning of desired test data. Every parameter set has an interpretation, and its effects can be easily

and intuitively predicted. The generation process scales with the number of points to generate, expanding the hierarchy as more elements are generated.

The primary weaknesses of the generator, on the other hand, are the trailing divisions and lack of fine control over specific areas of the structure. As mentioned before the hierarchy will always display a degree of self-similarity, replicating the same general form both vertically and horizontally. These issues can be at least partially addressed using post-processing procedures similar to the reassignment process. In its current form, the generator is also limited to normal distributed multidimensional and uncorrelated real value data.

Several potential improvements and extensions to the generator are a possibility, as well as further research into some aspects of its function. Firstly, the self-similar (fractal) nature of the hierarchies suggests a potential for the generator to be better described using the language of fractals and especially L-Systems [27]. Describing the generation process in this way may offer a more fine control over the details of the hierarchical structure. Secondly, post-processing procedures could be used to cut down on the number of trailing divisions, leading to cleaner hierarchies being generated and giving the user finer control over the overall structure. Thirdly, the procedure can be extended with different kernels leading to different structures or generators operating on different types of data.

As for the data generated, there is no single set of parameters that can be deemed better than any other. As a tool, the generator must be given a set of parameters considered based on the task the data will help test, model or develop. While the presented parameter sets can be used as a start, the parameters and their effects on the generated data should be considered and accordingly tweaked by the user.

References

References

[1] N. Sharma, A. Bajpai, R. Litoriya, Comparison the various clustering algorithms of weka tools, International Journal of Emerging Technology and Advanced Engineering 2 (2012) 73–80.

[2] M. Dash, H. Liu, P. Scheuermann, K. L. Tan, Fast hierarchical clustering and its validation, Data and Knowledge Engineering 44 (1) (2003) 109 – 138.

[3] R. Xu, D. Wunsch, II, Survey of clustering algorithms, Trans. Neur. Netw. 16 (3) (2005) 645–678. doi:10.1109/TNN.2005.845141.

[4] S. Dreiseitl, L. Ohno-Machado, H. Kittler, S. Vinterbo, H. Billhardt, M. Binder, A comparison of machine learning methods for the diagnosis of pigmented skin lesions, Journal of Biomedical Informatics 34 (1) (2001) 28 – 36.

[5] G. F. Cooper, C. F. Aliferis, R. Ambrosino, J. Aronis, B. G. Buchanan, R. Caruana, M. J. Fine, C. Glymour, G. Gordon, B. H. Hanusa, J. E. Janosky, C. Meek, T. Mitchell, T. Richardson, P. Spirtes, An evaluation of machine-learning methods for predicting pneumonia mortality, Artificial Intelligence in Medicine 9 (2) (1997) 107 – 138.

[6] P. Douglas, S. Harris, A. Yuille, M. S. Cohen, Performance comparison of machine learning algorithms and number of independent components used in fmri decoding of belief vs. disbelief, NeuroImage 56 (2) (2011) 544 – 553, multivariate Decoding and Brain Reading.

[7] C. Kampichler, R. Wieland, S. Calmé, H. Weissenberger, S. Arriaga-Weiss, Classification in conservation biology: A comparison of five machine-learning methods, Ecological Informatics 5 (6) (2010) 441 – 450.

[8] M. Meila, D. Heckerman, An experimental comparison of several clustering and initialization methods, CoRR abs/1301.7401.

[9] M. Steinbach, G. Karypis, V. Kumar, A comparison of document clustering techniques, in: KDD Workshop on Text Mining, 2000, pp. 525–526.

[10] G. Karypis, et al., Chameleon: A hierarchical clustering algorithm using dynamic modeling (1999).

[11] S. Guha, R. Rastogi, K. Shim, CURE: an efficient clustering algorithm for large databases, 1998, pp. 73–84.
URL http://citeseer.ist.psu.edu/guha98cure.html

[12] G. Costa, G. Manco, R. Ortale, E. Ritacco, Hierarchical clustering of {XML} documents focused on structural components, Data and Knowledge Engineering 84 (2013) 26 – 46. doi:http://dx.doi.org/10.1016/j.datak.2012.12.002.
URL http://www.sciencedirect.com/science/article/pii/S0169023X12001024

[13] U. Maulik, S. Bandyopadhyay, Performance evaluation of some clustering algorithms and validity indices, IEEE Trans. Pattern Anal. Mach. Intell. 24 (12) (2002) 1650–1654. doi:10.1109/TPAMI.2002.1114856.
URL http://dx.doi.org/10.1109/TPAMI.2002.1114856

[14] A. Mirzaei, M. Rahmati, M. Ahmadi, A new method for hierarchical clustering combination, Intell. Data Anal. 12 (6) (2008) 549–571.

[15] P. Cimiano, A. Hotho, S. Staab, Comparing conceptual, divise and agglomerative clustering for learning taxonomies from text, in: Proc. of the 16th Eureopean Conf. on AI, Spain, 2004, 2004, pp. 435–439.

[16] T. S. Madhulatha, An overview on clustering methods, CoRR abs/1205.1117.

[17] S. Lessmann, B. Baesens, C. Mues, S. Pietsch, Benchmarking classification models for software defect prediction: A proposed framework and novel findings, IEEE Trans. Softw. Eng. 34 (4) (2008) 485–496. doi:10.1109/TSE.2008.35.
URL http://dx.doi.org/10.1109/TSE.2008.35

[18] F. Murtagh, P. Contreras, Methods of hierarchical clustering, CoRR abs/1105.0121.
URL http://arxiv.org/abs/1105.0121

[19] Z. Ghahramani, M. I. Jordan, R. P. Adams, Tree-structured stick breaking for hierarchical data, in: NIPS, 2010, pp. 19–27.

[20] M. Spytkowski, H. Kwasnicka, Hierarchical clustering through bayesian inference., in: N. T. Nguyen, K. Hoang, P. Jedrzejowicz (Eds.), ICCCI (1), Vol. 7653 of Lecture Notes in Computer Science, Springer, 2012, pp. 515–524.

[21] Ł. P. Olech, M. Paradowski, Hierarchical Gaussian Mixture Model with Objects Attached to Terminal and Non-terminal Dendrogram Nodes, Springer International Publishing, Cham, 2016, pp. 191–201. doi:10.1007/978-3-319-26227-7_18.
URL http://dx.doi.org/10.1007/978-3-319-26227-7_18

[22] C. Blundell, Y. W. Teh, K. A. Heller, Bayesian rose trees, CoRR abs/1203.3468.
URL http://arxiv.org/abs/1203.3468

[23] M. Spytkowski, Ł. P. Olech, H. Kwaśnicka, Hierarchy of Groups Evaluation Using Different F-Score Variants, Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 654–664. doi:10.1007/978-3-662-49381-6_63.
URL http://dx.doi.org/10.1007/978-3-662-49381-6_63

[24] A. Shiravi, H. Shiravi, M. Tavallaee, A. A. Ghorbani, Toward developing a systematic approach to generate benchmark datasets for

intrusion detection, Computers and Security 31 (3) (2012) 357 – 374. doi:http://dx.doi.org/10.1016/j.cose.2011.12.012.
URL http://www.sciencedirect.com/science/article/pii/S0167404811001672

[25] J. Chen, Comparison of clustering algorithms and its application to document clustering, Ph.D. thesis, Princeton University, The Department of Computer Science, Princeton, NJ, USA, aAI3169789 (2005).

[26] J. O'shea, Z. Bandar, K. Crockett, A new benchmark dataset with production methodology for short text semantic similarity algorithms, ACM Trans. Speech Lang. Process. 10 (4) (2014) 19:1–19:63.

[27] P. Prusinkiewicz, J. Hanan, Lindenmayer systems, fractals, and plants, Vol. 79, Springer Science & Business Media, 2013.